

Mobile Ad Hoc Networking and the IETF

Joseph P. Macker^a

macker@itd.nrl.navy.mil

^a Information Technology Division, Naval Research Laboratory Washington, DC, USA

M. Scott Corson^b

corson@isr.umd.edu

^b Institute for Systems Research, University of Maryland, College Park, MD, USA

This article is the first in a series intended to chronicle the status of work underway within the Mobile Ad hoc NETWORKS (MANET) Working Group of the Internet Engineering Task Force (IETF). This article provides a short history and high-level, conceptual tutorial of MANET technology. We present an overview of the working group's vision and charter, and a glimpse into a technical architecture under consideration for achieving this vision.

I. Introduction

Recent advances in portable computing and wireless technology are opening up exciting possibilities for the future of mobile networking. The vision of the nascent Internet Engineering Task Force (IETF) Mobile Ad hoc Networking effort is helping to promote this future by providing improved standardized routing functionality to support self-organizing, mobile networking infrastructures. Applications of mobile ad hoc networking technology include industrial, commercial, and military communication networks involving cooperative mobile data exchange where wireless mobile nodes comprise the communications infrastructure.

Since its inception nearly thirty years ago, the Internet has existed as a network with a fundamentally *quasi-static* network infrastructure or topology¹. Its topology (see Fig. 1) consisted principally of Local Area Networks (LANs), typically referred to as subnetworks or "subnets", interconnected via gateways or routers into a larger "network of networks". While the Internet was architected in a decentralized fashion, and designed to adapt to topological changes (due perhaps to link outages, router failures, etc.) via dynamic routing, its routing technology was not designed for mobile networks where frequent topological changes may be the norm.

For a variety of reasons such as the advent of laptop-based mobile computing, the increasing use of Wireless LAN (WLAN) technology, and the integration of Internet Protocol (IP) data forwarding into cellular and microcellular systems, it has become desirable to be able to roam anywhere within the Internet and dynamically attach to its fixed router infrastructure in a fashion that gives *location transparency* to the mobile or nomadic user. Mobile IP was designed to support this type of nomadicity or "roaming", and does so with a combination of location management, tunneling and security mechanisms. Mobile ad hoc networking is solving a different problem by extending IP connectivity into the realm of autonomous, mobile, wireless domains, where the mobile nodes themselves form the network routing infrastructure or "fabric" in an ad hoc fashion.

Mobile ad hoc networking, while solving a separate problem than Mobile IP, shares a common goal—viz. extending




Figure 1: The Internet—a "network of networks"—where routers interconnect subnetworks.

seamless IP connectivity to mobile users. In this article, we present a brief history of the technology and give a high-level, conceptual tutorial. We then give the motivation and rationale for the formation of an IETF working group, give the current status of its work and the technical issues under consideration within the group.

II. Mobile Ad Hoc Networking

Mobile Ad Hoc Networking (a.k.a. Mobile Packet Radio Networking) is a name currently being given to a technology under development for the past 20 or so years, principally through research funding sponsored by the U.S. Government. Its initial sponsors included the Defense Advanced Research Projects Agency (DARPA), the U.S. Army and the Office of Naval Research (ONR). Significant early packet radio programs included the Survivable, Adaptive Networks (SURAN) Program, the Low-cost Packet Radio (LCR) Program and the Survivable Communication Networks (SCN) Program [1]. Today, government-sponsored work is still underway in networking programs such as the Tactical Internet and the Near-Term Digital Radio (NTDR).

A. The Technology

A Mobile Ad hoc NETWORK (MANET) consists of mobile platforms (each platform logically consisting of a router, possibly

¹Topology is a branch of mathematics concerned with those properties of geometric configurations (such as point sets) which are unaltered by elastic deformations (such as stretching) that are homeomorphisms. A network can be viewed abstractly as a graph whose "topology" at any point in time is defined by set of points (nodes) connected by edges (links).

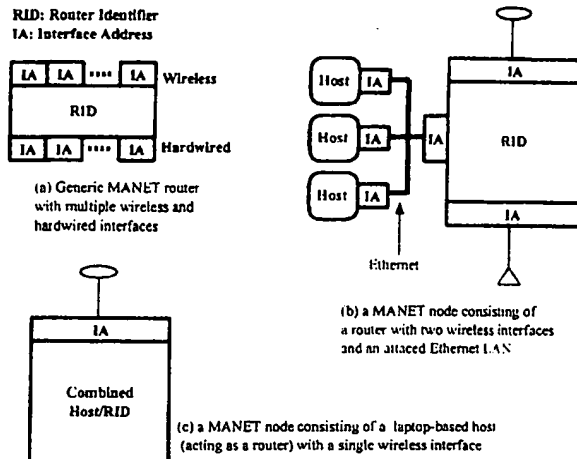


Figure 2: The generic MANET router structure and two possible MANET node configurations.

with multiple hosts and wireless communications devices, see Fig. 2)—herein simply referred to as “nodes”—which are free to move about arbitrarily. A MANET is an autonomous system of mobile nodes. The nodes may consist of separate, networked devices (see Fig. 2b), or may be integrated into a single device such as a laptop computer (see Fig. 2c). The nodes may be located in or on airplanes, ships, trucks, cars, perhaps even on people, and there may be multiple hosts per router.

The nodes are equipped with wireless transmitters and receivers using antennas which may be omnidirectional (broadcast), highly-directional (point-to-point) or some combination thereof. At a given point in time, depending on the nodes’ positions and their transmitter and receiver coverage patterns, transmission power levels and cochannel interference levels, a wireless connectivity in the form of a random, multihop graph or “ad hoc” network exists between the nodes (see Fig. 3).

This is in contrast with the topology of the existing Internet (recall Fig. 1), where the router topology is essentially static (barring network reconfiguration or router failures). In a MANET, the routers are *mobile* and interrouter connectivity may change frequently during normal operation. A MANET may operate either in isolation, or may be connected to the greater Internet via gateway routers (see Fig. 3).

MANETs have several salient characteristics:

1. **Dynamic topologies:** Nodes are free to move arbitrarily; thus, the network topology—which is typically multihop—may change randomly and rapidly at unpredictable times. Adjustment of transmission and reception parameters such as power may also impact the topology.
2. **Bandwidth-constrained, variable capacity links:** Wireless links will continue to have significantly lower capacity than their hardwired counterparts. One effect of the relatively low to moderate link capacities is that *congestion* is typically the norm rather than the exception, i.e. aggregate application demand will likely approach or exceed network capacity frequently.

3. **Power-constrained operation:** Some or all of the nodes in a MANET may rely on batteries for their energy. For these nodes, the most important system design criteria for optimization may be that of power conservation.

4. **Limited physical security:** Mobile wireless networks are generally more prone to physical security threats than are fixed, hardwired nets. Existing link security techniques are often applied within wireless networks to reduce security threats.

In addition, some envisioned networks (e.g. mobile military networks or highway networks) may be relatively *large* (e.g. tens or hundreds of nodes per autonomous system). A need for good routing scalability is not unique to MANETs. However, in light of the preceding mobile characteristics, the mechanisms required to achieve scalability likely are. These characteristics create a set of underlying assumptions and performance concerns for protocol design which differ from those guiding the design of routing and other network control protocols within the higher-speed, quasi-static topology of the fixed Internet.




Figure 3: The Internet with MANETs: shown with MANETs both connected to the Internet and operating in isolation.

B. Potential Applications

Some applications of MANET technology could include industrial and commercial applications involving *cooperative* mobile data exchange, or simply allowing continued mobile access to the greater Internet through a MANET. In addition, self-organizing mobile networks can be operated as robust, inexpensive alternatives or enhancements to cell-based mobile network infrastructures. There are also existing and future military networking requirements for robust, IP-compliant data services within mobile wireless communication networks, and many of these networks consist of highly-dynamic autonomous topology segments. Also, the developing technologies of “wearable” computing and communications and other

such "micronetworking" technologies [2] may provide interesting applications for MANET technology. MANET technology can provide an extremely flexible method for establishing communications for fire/safety/rescue operations or disaster recovery scenarios requiring rapidly-deployable communications with survivable, efficient dynamic networking. There are likely additional applications for MANET technology which are not presently envisioned. It is, simply put, efficient IP-based routing technology for highly-dynamic, self-organizing wireless networks.

III. MANET Technology and IP Routing

Wireless networking poses many technical challenges at the link and physical layers for designers. Much work has been done and is ongoing in the fields of multiple access, waveform/coding design, QoS and/or priority scheduling schemes. This is important technical research and work is likely to continue for years to come.

In addition to lower protocol layers, there is a set of challenges presented by wireless mobility to the performance of higher layers (e.g., network, transport, application). Within the Internet Protocol Suite, *the internetwork protocol and its associated routing protocols are responsible for glueing disparate media and end systems together*. Improved internetwork layer routing performance is therefore desirable in mobile networking environments where there is little or no underlying fixed infrastructure, and where both routers and hosts are "on the move".

A. Traditional Design

Traditionally, mobile packet radio systems have been "stovepipe" systems using proprietary, vertically-integrated technology at all levels of network control. This was due, in part, to the need to extract maximum performance from relatively low capacity, yet high-cost system components. Such networks were typically characterized by the use of a *single* wireless technology—whose wireless connectivity formed a single wireless topology—with multiple access, routing and other network control protocols specifically tailored for that wireless or link-layer technology.

Recently, the continuing advances in computing and communications technologies are yielding relatively high-performance, yet low-cost computing and communication devices. In coming years, communication devices utilizing technologies such as spread-spectrum and impulse radio communications will become less expensive. In addition, it may become commercially feasible to develop multimode radios and communication devices (e.g. integrated personal digital assistants and/or cellular phones) which use multiple wireless technologies as well. This is already realizable with laptop computer technology.

B. IP-based Design Approach

These hardware advancements, coupled with the increasing use of IP technology in both commercial and military systems,

are resulting in a shift in design philosophy from closed, proprietary systems based on stovepipe solutions to Internet compatible standards-based systems.

The rationale is multifold—some specific reasons are:

- **Routing Flexibility, Efficiency and Robustness:** When multiple wireless technologies are available in a given mobile network (see Fig. 4), it is desirable that routing occur at the IP layer. The figure gives an example network consisting of mobile nodes (e.g. each could be a car or tank), where each node consists of a mobile router with two different wireless devices attached, as well as an attached Ethernet LAN containing several IP-addressable hosts and other devices. In general, the wireless connectivity and, hence, the network topology corresponding to each wireless technology will be different. Thus, adjacent nodes may be connected by one or both technologies. By routing at the IP layer, it is possible to flexibly, efficiently and robustly forward a packet through the wireless "fabric" consisting of the logical *union* of the topologies of the individual wireless technologies.

For example, in Fig. 4, a packet may initially be routed via wireless technology *A* for several hops, and then switched to technology *B* on subsequent hops because either more capacity is available there, or because no connectivity exists in technology *A*'s topology. In single-technology (i.e. "subnet-based") routing, lack of connectivity in topology *A* would either have caused the packet to be dropped, or its restriction to the slower technology *A* would have resulted in higher end-to-end latency. Thus, it can be seen that the ability to dynamically route between wireless technologies gives added flexibility to the routing algorithm including more robustness to topological changes and potentially higher performance as well.

- **Hardware Economies of Scale:** As wireless hardware becomes a commodity—increasingly performant and inexpensive—the open systems design approach maintains that only the Multiple Access (MAC) layer need directly reflect the characteristics of a given physical layer technology. While it is true that tightly-coupled routing and MAC layer design for wireless, multihop networks is generally most efficient, it is not clear that a slightly looser coupling between a standardized routing algorithm and a MAC layer—through a standardized interface definition—cannot achieve nearly the same level of performance at less cost.

It is desirable to have these interface definitions as common as possible to ease widespread deployment and heterogeneous operation. With this approach, a standardized method for routing at the IP layer that can be used on top of any wireless technology. This frees wireless hardware vendors from the burden of incorporating routing functionality into their products and permits them to focus on building communications hardware. Such devices need only be attached to IP routers; each addressable as an IP interface of the router. Sufficient information regarding the link-layer can be made available to the network layer




Figure 4: A MANET consisting of two wireless technologies (*A* and *B*), and their logical union which forms the “wireless fabric” for routing at the IP-layer. It is interesting to note here that the topology of a MANET resembles that of the larger Internet—only in microcosm; i.e. each mobile node—with its collection of hosts—resembles a subnet, and the routers route information between these “mobile subnets” through the wireless fabric.

via a standardized interface for improved performance whenever possible.

A mobile wireless routing fabric may be made up of many different types of wireless links and technologies. This is not required to make use of a MANET approach but as mentioned earlier, if multiple technologies are available, IP-layer routing may even increase performance potential. Such a technical architecture would facilitate mass manufacture of inexpensive wireless devices which could interoperate with each other directly via the link-layer, or indirectly via the IP layer, with the IP-layer routing providing the glue that binds the mobile fabric together.

- **Future Quality of Service (QoS) Support:** The characteristics of various wireless technologies will likely be different (e.g. differing capacities, multiple access techniques, support for QoS, etc.) and, depending on QoS traffic characteristics, it may be favorable to route certain traffic classes over specific technologies when possible, only resorting to other technologies, if feasible, when necessary. In these cases, IP-layer routing permits route selection or forwarding policies not possible when routing is constrained to operating within a single wireless media type, and facilitates integration with IP QoS mechanisms developed for the fixed Internet. The future of QoS-capable mobile routing decisions remains largely a

research question, but a MANET architecture will leave open the possibility of such support for future consideration.

IV. Current Status and Issues in the MANET Working Group

The recently-formed MANET Working Group (WG) [3] in the IETF's Routing Area is chartered to provide improved mobile routing and interface definition standards for usage within the Internet Protocol Suite. In so doing, it hopes to lay the foundation for an open, flexible and extensible architecture for MANET technology. This is a challenging task as there are many issues that must be balanced in these complex systems. This section delves into some of the issues facing the MANET WG.

Summarizing the preceding section, an internetwork layer routing solution is important for the following reasons:

- end user and application pressure for seamless internetworking will continue regardless of the underlying infrastructure (fixed or mobile);
- the “physical media independence” features of the IP layer are important to support mobile routing through heterogeneous wireless fabrics;
- connectionless datagram forwarding is a robust, sensible technical approach for mobile networking;
- definition of some common routing approaches and interface definitions provides future flexibility, and also improves the cost effectiveness of deployed systems.

A. Goal of Working Group

The near-term goal of the MANET WG is to standardize an intra-domain unicast routing protocol which provides one or more modes of operation, each mode specialized for efficient operation in a given mobile networking “context”, where a context is a predefined set of network characteristics. The standardization of new routing technology is necessary because existing IP routing protocols were never designed to work in a mobile network (a different set of design trade-offs were made during their development that were appropriate for the quasi-static, high-speed networking environment then envisioned). Numerous studies and simulations have demonstrated the limitations of such approaches when applied to mobile wireless networks. New protocols are emerging from the research community which can be used to provide significant routing improvements for mobile usage.

The working group will also address issues pertaining to security and interaction/interface with link-layer protocols and Internet security protocols. In the longer term, the group may look at the issues of layering more advanced mobility services on top of the initial routing technology it standardizes. These longer term issues will likely include investigating multicast and QoS extensions.

B. Routing Performance Issues

To operate efficiently in a mobile networking context, a protocol should be designed and deployed with an expected networking context firmly in mind. Judging the merit of a routing protocol's design requires *metrics*—both qualitative and quantitative—with which to scope and measure its suitability and performance. These metrics should be *independent* of any given routing protocol.

The following is a list of desirable qualitative properties:

1. **Distributed operation:** This is an essential property, but it should be stated nonetheless.
2. **Loop-freedom:** Not required per se in light of certain quantitative measures (performance criteria), but generally desirable to avoid worst-case phenomena, e.g. a small fraction of packets spinning around in the network for arbitrary time periods. Ad hoc solutions such as TTL values can bound the problem, but more structured and well-formed approaches are generally desirable and oftentimes lead to better overall performance.
3. **Demand-based operation:** Instead of assuming uniform traffic distribution within the network (and maintaining routing between all nodes at all times), let the routing algorithm adapt to the varying traffic pattern on a demand or as needed basis. If this is done intelligently, it will utilize network resources more efficiently.
4. **"Sleep" period operation:** As a result of power conservation, or some other need to be inactive, some nodes of a MANET may stop transmitting and/or receiving (even receiving requires power) for arbitrary time periods. A routing protocol should be able to accommodate such sleep periods without overly adverse consequences. This property may require close coupling with the link-layer protocol through a standardized interface.

The following is a list of some quantitative metrics that are appropriate for assessing the performance of any routing protocol:

1. **End-to-end data throughput and delay:** Statistical measures of data routing performance (e.g., means, variances, distributions) are important. These are the measures of routing protocol effectiveness—how well it does its job—as measured from the *external* perspective of other protocols that make use of routing.
2. **Efficiency:** If data routing effectiveness is the external measure of a protocol's performance, efficiency is the *internal* measure of its effectiveness. To achieve a given level of data routing performance, two different protocols may expend differing amounts of overhead, depending on their internal efficiency. Protocol efficiency may or may not directly affect data routing performance. If control and data traffic must share the same channel, and the channel's capacity is limited, then excessive control traffic more severely impacts data throughput performance. In particular, it is useful to track two ratios that illuminate the *internal* efficiency of a protocol in doing its job (there may be others that the authors have not considered):

- **Average number of data bits transmitted per data bit delivered:** This can be thought of as a measure of the efficiency of delivering data within the network.
- **Average number of control bits transmitted per data bit delivered:** This measures the efficiency of the protocol in expending control overhead to deliver data packets. Note that this should include not only the bits in the routing control packets, but also the bits in the header of the data packets. In other words, *anything that is not data is control overhead*, and should be counted as control overhead required by the algorithm.

Similar to other routing work in fixed networks (e.g., multicast routing) it is unlikely that one routing protocol or mode for mobile ad hoc networking is the best approach for all networking contexts. Parameters that define a *networking context* and that should be considered during protocol design, simulation and comparison include:

1. **Network size:** Measured as the number of nodes.
2. **Network connectivity:** The average degree of a node (i.e. the average number of neighbors of a node).
3. **Topological rate of change:** The rate with which a network's topology is changing.
4. **Link capacity:** Effective link speed, measured in bits/second, after accounting for losses due to multiple access, coding, framing, etc.
5. **Fraction of unidirectional links:** How effectively does a protocol perform as a function of the presence of unidirectional links?
6. **Traffic patterns:** Different types of traffic distribution experienced within a network (e.g., (1) uniform: all nodes are equally likely receivers and sources providing equivalent network load, (2) non-uniform: certain routing nodes are sourcing and/or receiving more network traffic than others). How does a protocol behave across these different loading scenarios? Also what are the issues when mixtures of different connection types are considered, e.g., short-lived, transactional vs. long-lived, streaming.
7. **Mobility:** When, and under what circumstances, are temporal and spatial topological correlation relevant to the performance of a routing protocol? In these cases, what is the most appropriate model for simulating node mobility in a MANET?
8. **Fraction and frequency of sleeping nodes:** How does a protocol perform in the presence of sleeping and awakening nodes?

The preceding lists are not exhaustive, and merely give an indication of the number of dimensions that should be considered in the evaluation of protocols during a standardization process. These protocol evaluation issues highlight performance metrics that can help promote meaningful comparisons and assessments of protocol performance. These issues

differentiate MANETs from traditional, hardwired, multihop networks. In essence, the wireless MANET environment is one of scarcity rather than abundance, wherein bandwidth is relatively limited, and energy resources may be as well.

It should be recognized that a routing protocol tends to be well-suited for particular networking contexts, and less well-suited for others. In putting forth a description of a protocol for consideration with the MANET WG, both its *advantages* and *limitations* should be mentioned so that the appropriate networking context(s) for its usage can be identified. These attributes of a protocol can typically be expressed *qualitatively* e.g., whether the protocol can or cannot support shortest-path routing. Qualitative descriptions of this nature permit broad classification of protocols, and form a basis for more detailed *quantitative* assessments of protocol performance. The routing technology standardized by the group should ideally function effectively over a wide range of networking contexts—from small, collaborative, ad hoc conferencing groups to larger mobile, multihop networks.

C. Other Issues

Broader issues also await discussion by the group such as:

- What are the unique *address management* issues for MANETs, if any? Are there appropriate addressing practices that may improve the operation and impact of MANET systems within the Internet? (see Fig. 2a).
- What unique performance issues and interactions arise when operating with other existing protocol mechanisms (e.g., ICMP)?
- What should be included in the *standardized interfaces* with the link layer?
- Should a common *network-layer control protocol* be adopted which implements functionality required by most routing protocols?
- What are the various *problem scopes* and *networking contexts* that are sufficiently different as to require separate routing approaches?
- What *simulation framework* should the group adopt for protocol comparison?
- Should MANETs be configurable to carry *transit* data traffic, or should they be restricted to *stub* network operation, the latter configuration only permitting data traffic which originates or terminates within a MANET? This will likely depend upon each specific MANET approach presented and the development of appropriate exterior gateway approaches.

Future articles in this series will likely focus on such issues, and on the various routing protocols under consideration for standardization by the working group.

V. Conclusions

In summary, the networking opportunities for MANETs are intriguing and the engineering tradeoffs are many and challenging. We briefly presented a description of ongoing work and a

vision for the future integration of technology into the Internet. We believe that there is a need for standardized routing and interface solution(s) for mobile networking support, and are promoting such efforts along with other interested participants through the IETF MANET WG. The future holds the possibility for deploying inexpensive, IP internetworking compatible solutions to form self-organizing, wireless routing fabrics.

References

- [1] B. Leiner, editor. "Proceedings of the IEEE—Special Issue on Packet Radio Networks", volume 75, January 1987.
- [2] F. Bennett, D. Clarke, J. Evans, A. Hopper, A. Jones, and D. Leask. "Piconet: Embedded Mobile Networking". *IEEE Personal Communications*, 4(5), October 1997.
- [3] Internet Engineering Task Force. MANET WG Charter. <http://www.ietf.org/html.charters/manet-charter.html>.